arXiv:1610.04804v1 [stat.ML] 16 Oct 2016

Dynamic Stacked Generalization for Node Classification on Networks

Zhen Han* Alyson Wilson†

Abstract

We propose a novel stacked generalization (stacking) method as a dynamic ensemble technique using a pool of heterogeneous classifiers for node label classification on networks. The proposed method assigns component models a set of functional coefficients, which can vary smoothly with certain topological features of a node. Compared to the traditional stacking model, the proposed method can dynamically adjust the weights of individual models as we move across the graph and provide a more versatile and significantly more accurate stacking model for label prediction on a network. We demonstrate the benefits of the proposed model using both a simulation study and real data analysis.

1 Introduction

Network data and its relational structure have garnered tremendous attention in recent years. A network is composed of nodes and edges, where nodes represent interacting units and edges represent their relationships [7].

Node classification or node labeling on a network is a problem where we observe labels on a subset of nodes and aim to predict the labels for the rest [1]. There are various kinds of labels; for example, demographic labels such as age, gender, location, or social interests; labels such as political parties, research interests, or research affiliations. Collective inference estimates the labels of a set of related nodes simultaneously given a partially observed network by exploiting the relational auto-correlation of connected units [10], and it has been demonstrated effective in reducing classification error for many applications [3, 17, 11, 20]. Common collective inference methods are the Iterative Classification Algorithm (ICA) [17], Gibbs Sampling (Gibbs), and Relaxation Labeling (RL) [12, 14].

In many cases, multiple types of relationships can be observed in the same network. For example, in a citation network, an edge can mean two papers have the same author, or they are published in the same journal, or one paper cites another. We may also observe additional node-level information, such as the title and abstract of a paper, which can potentially help increase the label classification accuracy. When multiple relations are present on a network, one can merge all the relations and sum the weights of common links to perform a typical collective classification [13]. An alternative is to combine all of the information through an ensemble framework. Fürnkranz [6] introduced hyperlink ensembles for classifying hypertext documents, where he suggests first predicting the label of each hyperlink attached to a document and then combining these individual predictions using ensembles to make a final prediction for the label of the target document. A different approach was proposed by Heß and Kushmerick [9], where they suggest training separate classifiers for the local and relational attributes and then combining the local and relational classifiers through voting. A local classifier is trained using only node-level, or local, features; for example, title, abstract, or year of publication. A relational classifier infers a node's label by using relational features; for example, the labels of the connected neighbors. Cataltepe et al. discussed a similar ensemble approach [2], where they considered different voting methods, such as the weighted average, average, and maximum. Eldardiry and Neville [4] discussed an across-models collective classification method that formed ensembles of the estimates from multiple classifiers using a voting idea similar to collective inference to reduce variance.

The above literature focuses on combining multiple classifiers through some type of aggregation. Preisach and Schmidt-Thieme [18] proposed to use stacking instead of a simple voting as a more robust and powerful generalizing method to combine predictions made by local and relational classifiers. They suggest training each classifier independently and combining the predicted class probabilities from a pool of local and relational classifiers through stacking, which assigns constant weights to each classifier in a supervised fashion.

Stacked generalization (stacking) [22] is a technique for combining multiple classifiers, each of which has been individually trained for a specific classification task, to achieve greater overall predictive accuracy. The method first trains individual classifiers using cross-validation on the training data. The original training data is called level-0 data, and the learned models are called level-0 classifiers. The prediction outcomes from the level-0 models are pooled for the second-stage learning, where a meta-classifier is trained. The pooled classification outcomes are called level-1 data and the meta-classifier is called the level-1 generalizer.

Ting and Witten [21] showed that for the task of classification, the best practice is to use the predicted class probabilities generated by level-0 models to construct level-1 data.

*MaxPoint Interactive.

†Department of Statistics, North Carolina State University.

Essentially, stacking learns a meta-classifier that assigns a set of weights to the class predictions made by individual classifiers. The traditional stacking model assumes the weight of each classifier is constant from instance to instance, which does not hold in general for many relational classifiers on a network. For example, the weighted-vote relational neighbor (wvRN) classifier [13] infers a node's label by taking a weighted average of the class membership probabilities of its neighbors. One expects that its performance might be dependent on a node's topological characteristics in the graph; for example, the number of connected neighbors. On the other hand, local classifiers that are trained using only a node's local attributes are less dependent on its topological features. Consequently, when we combine local and relational classifiers, it is beneficial to have a set of weight functions instead of constant weights for each classifier. There has been some previous work on dynamically ensemble local and relational models [16, 23]. However, they impose parametric models on the weight functions that are not flexible to capture complex weighting functions.

In this paper, we develop a dynamic stacking framework using a generalized varying coefficient model, which allows the weights for each classifier to vary smoothly with a node's topological characteristics in a non-parametric way. We illustrate the benefits of incorporating a node's topological features into stacking. To the best of our knowledge, this is the first work that considers non-parametric functional weight stacking.

2 Background and Motivation

Network data can be represented by a graph $G = (V, E, Y)$ with vertices (nodes) V, edges (connections) $E = \{v_1, v_2\}, v_1, v_2 \in V$, and labels Y. The graph G is partitioned into two sets of vertices, V_{train} and V_{test}, with $V_{\text{train}} \cup V_{\text{test}} = V$ and $V_{\text{train}} \cap V_{\text{test}} = \emptyset$. We are given a classification problem with C classes. Class labels, y_i, are observed for nodes in the training set $v_i \in V_{\text{train}}$, while the labels of the test set V_{test} are unknown and need to be estimated. A relational classifier uses the attributes and/or labels from a node's connected neighbors to make predictions. However, unlike a typical classification problem, a node's neighbors may have missing attributes and/or labels, which in turn need to be estimated. Collective inference [11, 20] has been developed to make joint inference on the test nodes and produce consistent results.

We examinine the Cora [15] and the PubMed Diabetes [20] data sets, where we evaluate the collective classification accuracy on nodes with various topological characteristics. We consider the wvRN classifier as the relational classifier [13], defined as follows, and the Iterative Classification Algorithm (ICA) [12, 14] for collective inference as defined in Algorithm 1.

Algorithm 1 Iterative Classification Algorithm (ICA)

1: For $v_i \in V_{test}$, initialize the node labels, y_i, with a dummy label **null**.
2: **repeat**
3: Generate a random sequence of nodes, O, in V_{test}
4: **for** node $v_i \in O$ **do**
5: Apply the relational classifier model, using only non-**null** labels from $\mathbb{N}_i$, the neighborhood of v_i, and output an estimated class membership probability vector. We ignore nodes that have not been classified, so if all labels in $\mathbb{N}_i$ are **null**, we assign the label **null** to v_i.
6: Assign the label, c, with the largest class membership probability to v_i.
7: **end for**
8: **until** class assignments for V_{test} stop changing or a maximum number of iterations is reached.

DEFINITION 2.1. For a given node $v_i \in V_{\text{test}}$, the wvRN classifier estimates the class probability $P(y_i|\mathbb{N}_i)$ by the weighted average of the class membership probabilities in the neighborhood of v_i, $\mathbb{N}_i$:

$$P(y_i = c|\mathbb{N}_i) = \frac{1}{Z} \sum_{v_j \in \mathbb{N}_i} w_{i,j} P(y_j = c|\mathbb{N}_j), \tag{2.1}$$

where Z is a normalizing constant and $w_{i,j}$ is the weight associated with the edge between v_i and v_j.

Macskassy and Provost showed that the weighted-vote relational neighbor classifier is equivalent to the Gaussian-field model [14].

The Cora data set is a public academic database composed of papers from Computer Science. It contains a citation graph with attributes/labels of each paper (including authors, title, abstract, book title, and topic labels). We only consider the topics of each paper as its label and ignore the other attributes. We remove papers with no topic labels and construct the data set by keeping the largest connected component in the network. The final data set is an unweighed and non-directional network, with 19,355 nodes and 58,494 edges. Labels are 70 topic categories, and each paper is classified into one of the categories.

We randomly sample 80% of nodes from V into V_{test} and set their labels to **null**. We then make predictions using ICA on the nodes in V_{test} and calculate the classification accuracy for different levels of degrees and closeness centrality. We repeat this experiment 100 times and the results are displayed in Figures (1) and (2). In Figure (1), the classification accuracy of the wvRN classifier is dependent on the degree of v_i. As the count of a node's neighbor increases from 1 to 10, the average classification accuracy jumps from

45% to more than 60%. There are a limited number of nodes with degrees greater than 10, and thus the variance of the average classification accuracy goes up considerably. Closeness centrality is defined as the reciprocal of a node's total distance from all other nodes, which relates to the idea of "being in the middle of things." Unlike degree, closeness centrality is a continuous variable. We binned its range into 100 equal-length intervals and calculated classification accuracy in each bin. From Figure (2), we observe a steady upward trend in the classification accuracy near the center of the spectrum. There are not many nodes near the two ends of the spectrum, and this contributes to the large variation in the classification accuracy.

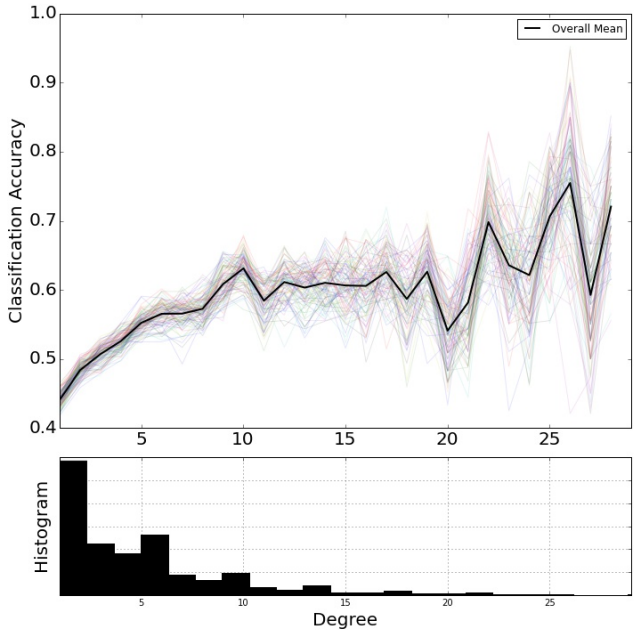

Figure 1: Classification accuracy for the relational classifier at different levels of node degree in the Cora data set.

We performed the same analysis on the Pubmed Diabetes data set. The Pubmed data set is a medical database composed of diabetes-related medical papers derived from the PubMed database. The graph is a citation network with 19,717 papers and 44,338 edges. Each publication is assigned one of three categories as its label and a TF/IDF weighted word vector as an extra attribute. Here we ignore the extra attributes and only consider the topic category labels. We observe results similar to those from the Cora data set in Figures (3) and (4).

3 Dynamic Stacking Model

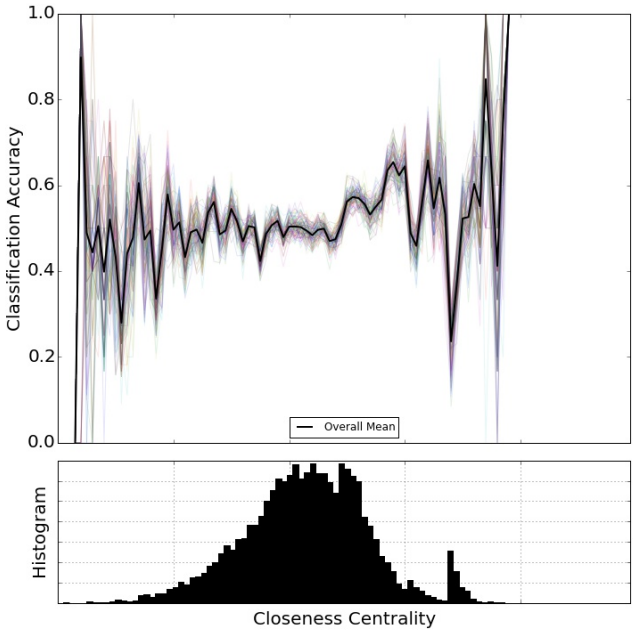

Figure 2: Classification accuracy for the relational classifier at different levels of node closeness centrality in the Cora data set.

3.1 Notation for the Stacked Generalization Model

Stacked generalization (stacking) is a general method for combining multiple lower-level models to improve overall predictive accuracy [22, 21]. Here we follow the notation in [21]. Given data $\mathbb{D} = \{(y_i, \boldsymbol{x}_i) \text{ for } i = 1, \cdots, N\}$, let $\boldsymbol{x}_i$ be the feature vector and y_i the label of the i-th observation. Here we focus on categorical responses for y, and assume y has C categories. We first randomly partition the data into J roughly equal-sized parts $\mathbb{D}_1, \mathbb{D}_2, \cdots, \mathbb{D}_J$. Define $\mathbb{D}_j$ and $\mathbb{D}_{-j} = \mathbb{D} - \mathbb{D}_j$ to be the test and training data sets for the j-th fold of a J-fold cross validation.

Suppose we have K classifiers. We train each of the K classifiers using the training set $\mathbb{D}_{-j}$ with results $\mathbb{M}_k$. $\mathbb{M}_1, \cdots, \mathbb{M}_K$ are called *level-0 models*. We then apply the K classifiers on the test set $\mathbb{D}_j$ and denote $\boldsymbol{z}_{ik} = \mathbb{M}_k(\boldsymbol{x}_i), \boldsymbol{x}_i \in \mathbb{D}_j$ as the estimated class probability vector from $\mathbb{M}_k$ for $\boldsymbol{x}_i \in D_j$. We repeat this process for $j = 1, \cdots, J$ and collect the outputs from the K models to form the *level-1 data* as follows:

$$D_{\text{level 1}} = \{(y_i, \boldsymbol{z}_{i1}, \cdots, \boldsymbol{z}_{iK}), \text{ for } i = 1 \cdots, N\}. \tag{3.2}$$

$\boldsymbol{z}_{ik}$ is the predicted class probability vector from $\mathbb{M}_K$ for observation i, and therefore $\sum_c z_{ikc} = 1$. We drop the last element z_{ikC} from vector $\boldsymbol{z}_{ik}$ to avoid multicollinearity issues. We then fit a supervised classification model, $\tilde{\mathbb{M}}$, using the level-1 data, which is called the *level-1 generalizer*.

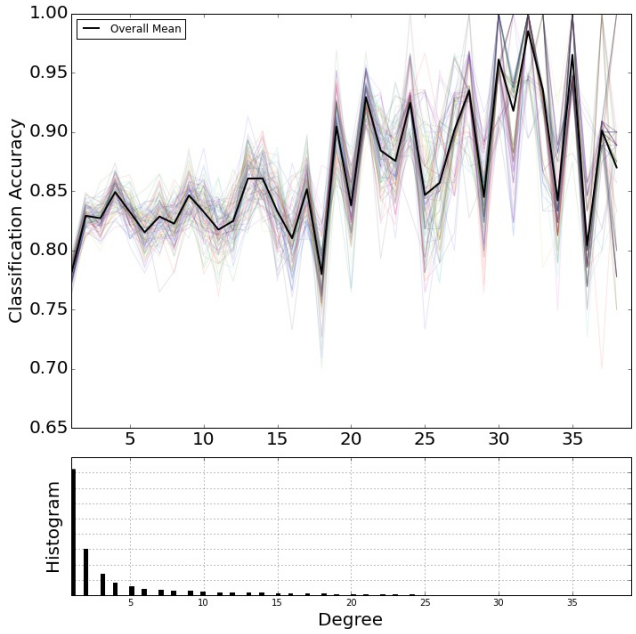

Figure 3: Classification accuracy for the relational classifier at different levels of node degree in the PubMed data set.

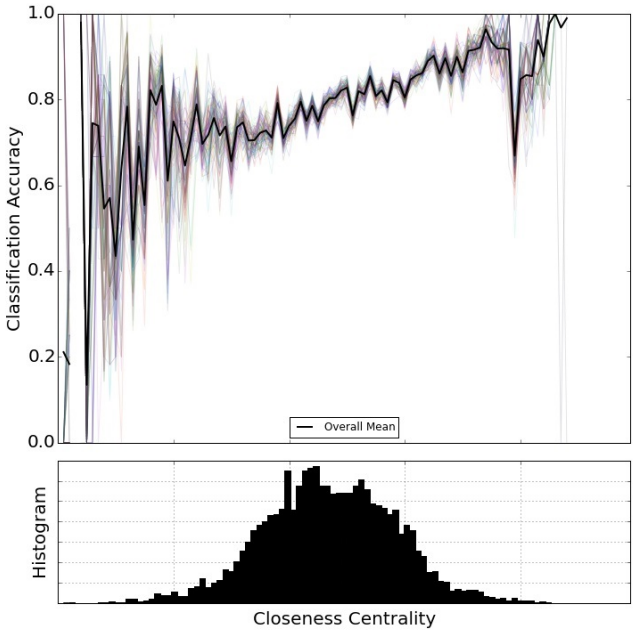

Figure 4: Classification accuracy for the relational classifier at different levels of node closeness centrality in the PubMed data set.

For prediction, a new observation $\boldsymbol{x}_{\text{new}}$ is input into the K low-level classifiers, $\mathbb{M}_1, \cdots, \mathbb{M}_K$. The estimated class probability vectors, $\boldsymbol{z}_1, \cdots, \boldsymbol{z}_K$, are then concatenated and input into $\tilde{\mathbb{M}}$, which outputs the final class estimate for that observation.

For classification on networks, Preisach and Schmidt-Thieme [18] adapted the stacking technique and combined a local classifier with a relational classifier using a logistic regression model as the level-1 generalizer. However, in their paper, the coefficients in the logistic regression are constant, meaning the weights of individual component classifier are "static" from node to node. From previous observations in Figures (1), (2), (3), and (4), the accuracy of a relational classifier is often dependent on some topological feature of a node. Therefore, it could be beneficial to "dynamically" allocate the weights of individual classifiers based on some other variable. We discuss a dynamic stacking model using a generalized varying coefficient model in the next section to account for this observation. Compared with competing methods that also consider dynamic stacking [16, 23], the proposed model is non-parametric, more flexible, and can learn more complex weighting functions.

3.2 Generalized Varying Coefficient Model through Smoothing Splines Here we develop a dynamic stacking model using a generalized varying coefficient model. Instead of having a set of constant coefficients in the regression, we allow the coefficients to vary as smooth functions of other variables. The generalized varying-coefficient model was proposed by Hastie and Tibshirani [8] and was reviewed in [5].

Similar to traditional stacked generalization, the inputs for the dynamic stacking model are the assembled outputs from multiple level-1 classifiers, along with an extra covariate: $\{(y_i, \boldsymbol{Z}_i, u_i),\ \text{for } i = 1, \cdots, N\}$. y_i is the true class label of an observation, $\boldsymbol{Z}_i$ is the concatenated predicted class membership vector from a pool of inhomogeneous classifiers with dimension p. Each of the component classifiers could potentially look at a different set of features of an instance and make a prediction from its point of view. u_i is an "extra" covariate of a observation, which presumably would affect the prediction accuracy made by at least one classifier. Here we focus on the case where y_i is binary and u_i is continuous. One can easily extend this method to multi-class classification problems by using a one-vs-all strategy.

The regression function is modeled as:

$$
\begin{aligned}
g(m(U_i, \boldsymbol{Z}_i)) =& \beta_0 + \boldsymbol{Z}_i^T \boldsymbol{\beta}(U_i) \\
=& \beta_0 + \sum_{j=1}^{p} Z_{ij} \beta_j(U_i)
\end{aligned} \tag{3.3}
$$

where $g(\cdot)$ is the logit link function, $\boldsymbol{\beta}(\cdot)$ is the functional coefficient vector that varies smoothly with an extra scalar

covariate, and β_0 is a constant intercept. Instead of a constant intercept, one can trivially add a functional intercept by appending 1 to Z_i. However, in this paper, we focus on the constant intercept case. We assume that each functional coefficient $\beta_j(\cdot)$ for $j = 1, \cdots, p$ can be approximated by spline functions:

$$\beta_j(\cdot) = \sum_{k=1}^{K_j} \eta_{jk} B_{jk}(\cdot), \text{ for } j = 1, \cdots, p, \tag{3.4}$$

where for each β_j, $B_{jk}(\cdot)$ for $k = 1, \cdots, K_j$ is a set of spline basis functions. Without loss of generality, we use the same set of B-spline basis functions for all $\beta_1(\cdot), \cdots, \beta_p(\cdot)$. Henceforth, we denote the set of B-spline basis functions as $B_1(\cdot), \cdots, B_K(\cdot)$, where K is the number of basis functions. We can then rewrite equation (3.4) as:

$$\beta_j(\cdot) = \sum_{k=1}^{K} \eta_{jk} B_k(\cdot) \text{ for } j = 1, \cdots, p. \tag{3.5}$$

We substitute equation (3.5) into equation (3.3) and rewrite the regression function as

$$\begin{aligned} g(m(U_i, \boldsymbol{Z_i})) &= \beta_0 + \sum_{j=1}^{p} \sum_{k=1}^{K} \boldsymbol{Z}_{ij} \eta_{jk} B_k(U_i) \\ &= \beta_0 + \boldsymbol{Z}_i^T \boldsymbol{B}(U_i) \boldsymbol{\eta}, \end{aligned} \tag{3.6}$$

Denote $\boldsymbol{B}_*(U) = (B_1(U), \cdots, B_K(U))_{1\times K}$. We can express $\boldsymbol{B}(U)$ as:

$$\boldsymbol{B}(U) = \begin{bmatrix} \boldsymbol{B}_*(U) & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \boldsymbol{B}_*(U) \end{bmatrix}_{p \times pK}$$

and express $\boldsymbol{\eta}$ as:

$$\boldsymbol{\eta} = \begin{bmatrix} \eta_{11} \\ \vdots \\ \eta_{pK} \end{bmatrix}_{pK}$$

.

We can estimate $\beta_0, \beta_1(\cdot), \cdots, \beta_p(\cdot)$ by directly minimizing:

(3.7)

$$\begin{aligned} \hat{\beta}_0, \hat{\beta}_1(\cdot), \cdots, \hat{\beta}_p(\cdot) = & \underset{\beta_0, \beta_1(\cdot), \cdots, \beta_p(\cdot)}{\arg\min} \\ & - \sum_{i=1}^{N} \ell(\beta_0 + \sum_{j=1}^{p} Z_{ij} \beta_j(U_i), y_i) \\ & + \lambda \sum_{j=1}^{p} \int (\beta_j^{''}(x))^2 dx \end{aligned}$$

where $\ell(g(m(U_i, \boldsymbol{Z}_i)), y_i)$ is the log-likelihood function of the logistic regression, $\lambda \sum_{j=1}^{p} \int (\beta_j^{''}(x))^2 dx$ is a smoothness penalty term that controls the total curvature of the fitted $\beta_j(\cdot)$ for $j = 1, \cdots, p$, and λ is a a smoothing parameter that controls the trade-off between model fit and the roughness of the fitted $\beta_j(\cdot)$s. When $\lambda \to 0$, we have a set of wiggly $\beta_j(\cdot)$s; as $\lambda \to \infty$, the minimization of (3.7) will produce a set of linear $\beta_j(\cdot)$s.

For the constant intercept case, one can absorb β_0 into $\boldsymbol{\eta}$ as:

$$\boldsymbol{\eta}^* = \begin{bmatrix} \beta_0 \\ \eta_{11} \\ \dots \\ \eta_{pK} \end{bmatrix}_{1+pK}$$

and append a constant 1 to the beginning of the product, $\boldsymbol{Z}_i^T \boldsymbol{B}(U_i)$. We can write the optimization in equation (3.7) w.r.t. $\boldsymbol{\eta}^*$ as:

$$\hat{\boldsymbol{\eta}^*} = \underset{\boldsymbol{\eta}^*}{\arg\min} \{-\ell(\boldsymbol{\eta}^*) + \lambda \boldsymbol{\eta}^{*T} \boldsymbol{H} \boldsymbol{\eta}^*\}, \tag{3.8}$$

where $\boldsymbol{H}$ is the assembled penalty matrix:

$$\boldsymbol{H} = \begin{bmatrix} 0 & \dots\dots\dots & 0 \\ . & H^1 & 0 & 0 \\ . & \dots & H^j & \dots \\ 0 & 0 & 0 & H^p \end{bmatrix}_{(1+pK)\times(1+pK)}$$

H^j is the smoothness penalty matrix for $\beta_j(\cdot)$, and $\{H^j\}_{mn} = \int B_m^{''}(x) B_n^{''}(x) dx$ for $m, n = 1, \cdots, K$. Since we are using the same set of basis functions, $H^1 = \cdots = H^p$. It can be shown that $-\ell(\boldsymbol{\eta}^*)$ is convex w.r.t. $\boldsymbol{\eta}^*$. Also, one can show that $\boldsymbol{H}$ is positive semi-definite, so $\lambda \boldsymbol{\eta}^{*T} \boldsymbol{H} \boldsymbol{\eta}^*$ is convex w.r.t. $\boldsymbol{\eta}^*$ as well. Therefore, there exists a unique $\hat{\boldsymbol{\eta}^*}$ that optimizes equation (3.8).

Given a specified smoothness penalty parameter λ, to estimate $\boldsymbol{\eta}^*$, we employ an iterative Newton-type optimization method by directly calculating the derivatives of the objective function in equation (3.8). The smoothness penalty parameter λ can be chosen by cross-validation, where, for a range of λ values, we iteratively leave out a subset of the training data, fit the model using the rest of the data, and compute the prediction error on the held out data set. The best λ is set to the one with the smallest objective function value.

4 Simulation Study

Here we compare the performance of the dynamic stacking method against standard benchmarks using simulated data sets. [18] used a standard logistic regression model as the level-1 generalizer to combine a local and a relational classifier. [19] suggested that regularization is necessary to reduce over-fitting and increase predictive accuracy, and they considered lasso regression, ridge regression, and elastic net

regression. In our simulation study, we use lasso regression, ridge regression, and logistic regression as benchmark level-1 generalizers, and for each of the benchmark generalizers, we experiment with adding an additional covariate and/or interaction terms into the stacking, and compare their performance with the dynamic stacking model.

For the simulation, $N = 2000$ observations are generated for $i = 1, \cdots, N$. Z_{1i} and Z_{2i} are the predicted positive class probabilities from two classifiers, $\mathbb{Z}_1$ and $\mathbb{Z}_2$, and they are generated independently from a uniform distribution on $[0, 1]$. w_i is the error term which follows a normal distribution, $N(0, 1)$. u_i is an extra covariate which may affect the weight of a classifier, and it is generated from a uniform distribution on $[0, 1]$. Finally, the response y_i is generated from a Bernoulli distribution with $p(y_i = 1)$ specified as one of the following three cases. In case 1, the classifier weights are not dependent on u, while case 2 has linear dependence, and case 3 has non-linear dependence.

Case 1:

$$\text{logit}(p(y_i = 1)) = -3 + 3Z_{1i} + 3Z_{2i} + w_i$$

Case 2:

$$\text{logit}(p(y_i = 1)) = -3 + 3u_iZ_{1i} + 3Z_{2i} + w_i$$

Case 3:

$$\text{logit}(p(y_i = 1)) = -3 + 3\sin(6u_i)Z_{1i} + 3Z_{2i} + w_i$$

For training and evaluation, the N observations are evenly split into training and testing sets. We train the dynamic stacking model and benchmark methods using the training set, where the penalty parameters of the proposed method and benchmarks are selected by 10-fold cross-validation. The fitted models are then applied to predict on the testing set, and the final prediction accuracy on the test set is measured by the Area Under the Curve (AUC) from prediction scores as shown in Table (1). For methods[1] (Logistic[1], Lasso[1], and Ridge[1]), the inputs to the level-1 generalizer are $\{(y_i, Z_{1i}, Z_{2i})\}$. For methods[2], we add the additional covariate u into the input: $\{(y_i, Z_{1i}, Z_{2i}, u_i)\}$. For methods[3], in addition to u, we further add its interaction with Z_{1i}, Z_{2i} into the input: $\{(y_i, Z_{1i}, Z_{2i}, u_i, Z_{1i}u_i, Z_{2i}u_i)\}$.

From Table (1), the dynamic stacking model performs no worse than the standard methods under all scenarios. It has better performance than the "static" stacking models when there is an underlying non-linear dependency between a classifier's performance and the extra covariate. In case 1, the dynamic stacking model generalizes to the traditional stacking models and does not over fit the data. In case 2, where a linear dependency exists, the proposed model generalizes to methods[3], where interaction terms with the extra covariate are added into the "static" stacking model. In case 3, where a non-linear dependency exists, the dynamic stacking model outperforms all benchmarks.

Table 1: AUC score comparison between the proposed method and benchmarks. For level-1 generalizers, methods[1] use Z_{1i} and Z_{2i} as covariates, methods[2] contains U_i as an extra feature, and methods[3] further include linear interaction terms U_iZ_{1i}, and U_iZ_{2i}. The standard deviation of the accuracy score is calculated from 50 repetitions and is shown in parenthesis.

METHODS	CASE 1	CASE 2	CASE 3
RANDOM GUESS	0.49	0.50	0.51
	(0.02)	(0.02)	(0.02)
Z_1 ONLY	0.68	0.60	0.54
	(0.02)	(0.02)	(0.02)
Z_2 ONLY	0.68	0.68	0.67
	(0.02)	(0.02)	(0.02)
LOGISTIC[1]	0.75	0.71	0.67
	(0.02)	(0.02)	(0.01)
LASSO[1]	0.75	0.71	0.67
	(0.02)	(0.02)	(0.01)
RIDGE[1]	0.75	0.71	0.67
	(0.02)	(0.02)	(0.01)
LOGISTIC[2]	0.75	0.73	0.75
	(0.02)	(0.02)	(0.02)
LASSO[2]	0.75	0.72	0.74
	(0.02)	(0.02)	(0.02)
RIDGE[2]	0.75	0.72	0.74
	(0.02)	(0.02)	(0.02)
LOGISTIC[3]	0.75	0.73	0.76
	(0.01)	(0.02)	(0.01)
LASSO[3]	0.74	0.73	0.76
	(0.02)	(0.02)	(0.01)
RIDGE[3]	0.74	0.73	0.76
	(0.02)	(0.02)	(0.01)
PROPOSED METHOD	**0.75**	**0.73**	**0.79**
	(0.02)	**(0.02)**	**(0.01)**

5 Real Data Analysis

Here we revisit the Cora data set [15], where we use paper titles as node attributes and topic classification as labels. We remove nodes with no title or topic classifications, and the final graph contains 11,187 nodes and 33,777 edges. Seventy topic categories are used as labels, and each paper belongs to one of the categories. For simplicity, we convert the classification problem into a binary classification problem by giving a positive label if the topic falls under the `/Artificial_Intelligence/` category. We then use the closeness centrality of each node in the graph as an additional covariate in stacking.

We split all the nodes on the graph into a 20% training set and an 80% testing set. On the training set, we observe the titles and the topic classification label of each paper, while on the test set, we only observe the titles. We fit a local classifier using the word vector representation of its

Table 2: Mean accuracy score on 100 randomized experiments using Cora and PubMed data set.

	Xiang	McDowell	Lasso	Ridge	Logistic	Proposed
Cora	0.917476	0.918087	0.917898	0.918106	0.918401	**0.919270**
PubMed	0.879380	0.879205	0.889090	0.889125	0.889213	**0.891499**

Table 3: Pairwise accuracy score comparison between the proposed method versus each competing method in 100 randomized experiments. The values shown are the mean accuracy difference between the proposed method and each method and the p-value for H_0 : the proposed method is less accurate than that benchmark method.

Proposed	vs Xiang	vs McDowell	vs Lasso	vs Ridge	vs Logistic
Cora	0.1794	0.1183	0.1372	0.1164	0.0869
	p <0.01	**p <0.01**	**p <0.01**	**p <0.01**	**p <0.01**
PubMed	1.2119	1.2294	0.2409	0.2374	0.2286
	p <0.01	**p <0.01**	**p <0.01**	**p <0.01**	**p <0.01**

title only (Naive Bayes), and a relational classifier (ICA + wvRN) using only the labels from a paper's neighbor. We then fit a dynamic stacking model using the output from the two classifiers with their coefficients being smooth functions of the closeness centrality of a node. The smoothness penalty parameter is chosen by 10-fold cross-validation. One set of fitted coefficient curves for the two classifiers are shown in Figure 5. It allocates a higher weight on the relational classifier when a node has a high closeness centrality value and relies on the local classifier for nodes with a small closeness centrality value. This mirrors our observations from the previous discussion.

We compared the dynamic stacking model with the traditional stacking model on multiple standard level-1 generalizers (lasso, ridge, and logistic regression), all of which ignore the closeness centrality of a node during stacking. The penalty parameters for lasso and ridge regression are chosen by 10-fold cross-validation. We also implemented ensemble classification methods from [23, 16]. In [23], Xiang and Neville proposed a parametric weighting scheme to combine a local classifier and a relational classifier where model parameters are chosen by cross-validation. In [16], outputs from local and relational classifier are combined thorough a concept of label regularization and the model they proposed has no additional parameters. We repeat the train-test process 100 times randomly and record the accuracy score for each run.

We also performed model comparison using the PubMed data with the same general setup. Input to the local classifier is the TF/IDF representation of a paper, and we fit a dynamic stacking model using the output from the local classifier and the relational classifier with their coefficients being smooth functions of the degree centrality of a node. The classification accuracy comparison between the proposed method and the benchmarks is shown in Table (2) and

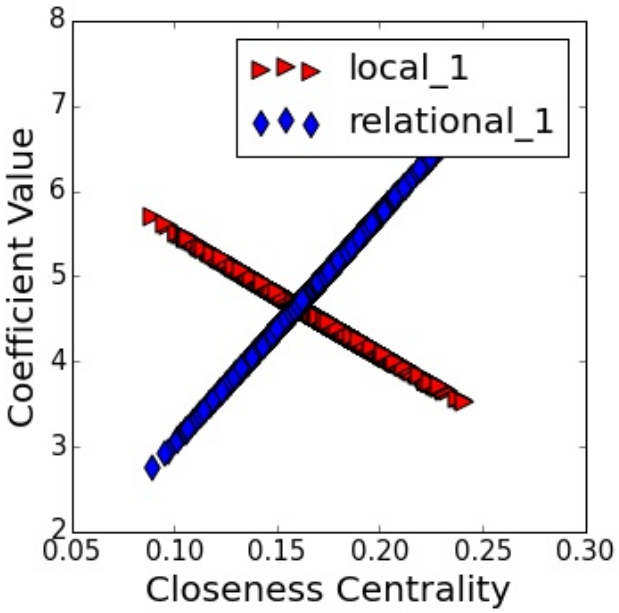

Figure 5: One set of fitted coefficient curves.

Table (3), where the accuracy is defined as $\sum_{i=1}^{N} \mathbb{I}_{y_i=\hat{y}_i}/N$ where $\hat{y}_i = \mathbb{I}_{\hat{p}_i>0.5}$.

By assuming the normality of the classification accuracy difference distribution, the dynamic stacking mode outperforms all benchmarks at p-value < 0.01. For the Cora data set, Figure 6 shows the source of the accuracy improvement. For each of the 100 repetitions, we calculate the difference in the absolute number of correctly classified nodes at different closeness centrality levels between the proposed model and benchmarks. The dynamic stacking model outperforms

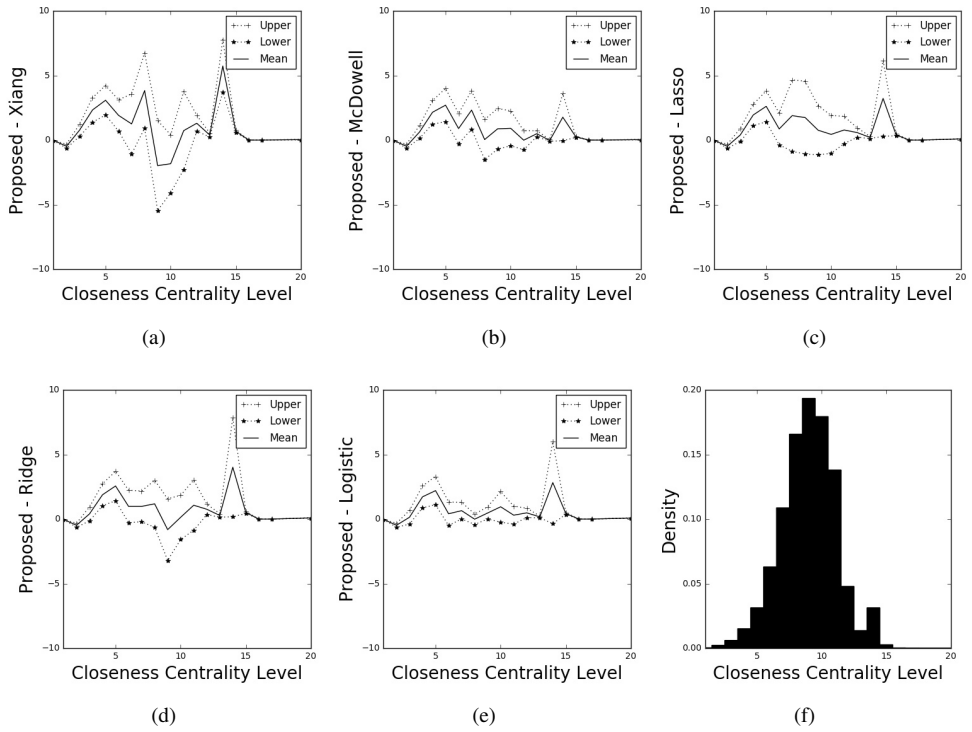

Figure 6: For each of the 100 repetitions, we calculate the difference in the number of correctly classified nodes at different closeness centrality levels between the dynamic stacking model and benchmarks. In (a) – (e), we calculate the group mean and a 95% confidence interval. (f) shows a density distribution of the closeness centrality for all nodes in the graph.

the benchmarks near the two ends of the closeness centrality spectrum where the balance between the local and relational classifier shifts considerably. For the majority of nodes in this data set, their closeness centrality clusters tightly around a specific value, which leaves little room for the dynamic stacking model to improve much beyond its static-weight counterparts in terms of the overall accuracy. However, for the nodes that are near the two extremes of the closeness centrality spectrum, we do see a significant improvement by using the dynamic stacking method.

6 Discussion

In this paper, by examining two public data sets, Cora and PubMed, we illustrate the motivation for incorporating node topological characteristics into stacked generalization for node classification on networks. We then develop a novel dynamic stacking method with functional weights for component models, each of which can vary smoothly with an extra covariate. Simulation studies show that the proposed method generalizes well to the benchmarks when the data does not present complex patterns, and outperforms all benchmarks otherwise. Real data analysis using Cora and PubMed shows the proposed method has a small yet significant improvement on the classification accuracy compared with traditional stacking models. Further analysis shows that most of the accuracy improvement comes from nodes near the two extremes of the closeness centrality spectrum where the balance between the local and relational classifier shifts the most. The limited number of nodes in that region explains the small improvement on the overall classification accuracy. However, for the nodes near the extremes of closeness centrality, we do see a considerable improvement on the classification accuracy using the dynamic stacking method. Overall, the dynamic stacking model allows the composition of the stacking model to change as we move across the network, and thus it potentially provides a more versatile and accurate stacking model for label prediction on a network.

7 Future Work

The proposed dynamic stacking model is a direct extension of the traditional stacking model, which uses logistic regression as level-1 generalizer. As discussed in [19], this model tends to overfit, especially when combining a large pool of noisy classifiers. To mitigate this problem, one can add a group-lasso penalty over the model coefficients, $\boldsymbol{\eta}^*$, into equation (3.8). Coefficients can be naturally

grouped if they are the basis coefficients for the same $\beta_j(\cdot)$: $\{(\eta_{j1}, \cdots, \eta_{jK}),$ for $j = 1, \cdots, p\}$. By adding a group-lasso penalty, the dynamic stacking model is more robust to the noise in the level-1 generalizing process.

8 Acknowledgment

This material is based upon work supported in part with funding from the Laboratory for Analytic Sciences (LAS). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the LAS and/or any agency or entity of the United States Government.

References

[1] Smriti Bhagat, Graham Cormode, and S. Muthukrishnan. Node classification in social networks. *CoRR*, abs/1101.3291, 2011.

[2] Zehra Cataltepe, Abdullah Sonmez, Kadriye Baglioglu, and Ayse Erzan. Collective classification using heterogeneous classifiers. In Petra Perner, editor, *MLDM*, volume 6871 of *Lecture Notes in Computer Science*, pages 155–169. Springer, 2011.

[3] Soumen Chakrabarti, Byron Dom, and Piotr Indyk. Enhanced hypertext categorization using hyperlinks. In *Proceedings of the 1998 ACM SIGMOD International Conference on Management of Data*, SIGMOD '98, pages 307–318, New York, NY, USA, 1998. ACM.

[4] Hoda Eldardiry and Jennifer Neville. Across-model collective ensemble classification. In *in AAAI*, 2011.

[5] J. Fan and W. Zhang. Statistical methods with varying coefficient models. *Statistics and its Interface*, 1:179–195, 2008.

[6] Johannes Fürnkranz. Hyperlink ensembles: A case study in hypertext classification. *Information Fusion*, 3:299–312, 2001.

[7] Anna Goldenberg, Alice X. Zheng, Stephen E. Fienberg, and Edoardo M. Airoldi. A survey of statistical network models. *Found. Trends Mach. Learn.*, 2(2):129–233, February 2010.

[8] Trevor Hastie and Robert Tibshirani. Varying-Coefficient Models (with discussion). *Journal of the Royal Statistical Society. Series B (Methodological)*, 55(4):757–796, 1993.

[9] Andreas Heß and Nick Kushmerick. Iterative ensemble classification for relational data: A case study of semantic web services. In *In Proceedings of the 15th European Conference on Machine Learning*. Springer, 2004.

[10] David Jensen, J. Neville, and Jennifer Neville. Linkage and autocorrelation cause feature selection bias in relational learning. In *In Proceedings of the 19th International Conference on Machine Learning*, pages 259–266. Morgan Kaufmann, 2002.

[11] David Jensen, Jennifer Neville, and Brian Gallagher. Why collective inference improves relational classification. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Seattle, Washington, USA, August 22-25, 2004*, pages 593–598, 2004.

[12] Qing Lu and Lise Getoor. Link-based classification. In Tom Fawcett and Nina Mishra, editors, *ICML*, pages 496–503. AAAI Press, 2003.

[13] Sofus A. Macskassy and Foster Provost. A simple relational classifier. In *Proceedings of the Second Workshop on Multi-Relational Data Mining (MRDM-2003) at KDD-2003*, pages 64–76, 2003.

[14] Sofus A. Macskassy and Foster Provost. Classification in networked data: A toolkit and a univariate case study. *J. Mach. Learn. Res.*, 8:935–983, May 2007.

[15] Andrew K. McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3(2):127–163, 2000.

[16] Luke McDowell and David W. Aha. Semi-supervised collective classification via hybrid label regularization. In *ICML*. icml.cc / Omnipress, 2012.

[17] J. Neville and D. Jensen. Iterative classification in relational data. pages 13–20. AAAI Press, 2000.

[18] Christine Preisach and Lars Schmidt-Thieme. Ensembles of relational classifiers. *Knowl. Inf. Syst.*, 14(3):249–272, March 2008.

[19] Sam Reid and Greg Grudic. Regularized linear models in stacked generalization. In *Multiple Classifier Systems*, volume 5519 of *Lecture Notes in Computer Science*, pages 112–121. Springer Berlin Heidelberg, 2009.

[20] Prithviraj Sen, Galileo Mark Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.

[21] Kai Ming Ting and Ian H. Witten. Stacked generalization: when does it work? In *in Procs. International Joint Conference on Artificial Intelligence*, pages 866–871. Morgan Kaufmann, 1997.

[22] David H. Wolpert. Stacked generalization. *Neural Networks*, 5:241–259, 1992.

[23] Rongjing Xiang and Jennifer Neville. Understanding propagation error and its effect on collective classification. In *11th IEEE International Conference on Data Mining, ICDM 2011, Vancouver, BC, Canada, December 11-14, 2011*, pages 834–843, 2011.